



**US Army Corps
of Engineers**

Waterways Experiment
Station

Technical Report ITL-97-3
May 1997

U.S. Army Engineer Waterways Experiment Station Major Shared Resource Center Benchmark Results and Analysis; Performance Level 1

by John E. West, Alex R. Carrillo

DTIC QUALITY INSPECTED 3

Approved For Public Release; Distribution Is Unlimited

19970711 034

DTIC QUALITY INSPECTED 3

Prepared for Headquarters, U.S. Army Corps of Engineers

The contents of this report are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval of the use of such commercial products.

The findings of this report are not to be construed as an official Department of the Army position, unless so designated by other authorized documents.



PRINTED ON RECYCLED PAPER

U.S. Army Engineer Waterways Experiment Station Major Shared Resource Center Benchmark Results and Analysis; Performance Level 1

by John E. West, Alex R. Carrillo
U.S. Army Corps of Engineers
Waterways Experiment Station
3909 Halls Ferry Road
Vicksburg, MS 39180-6199

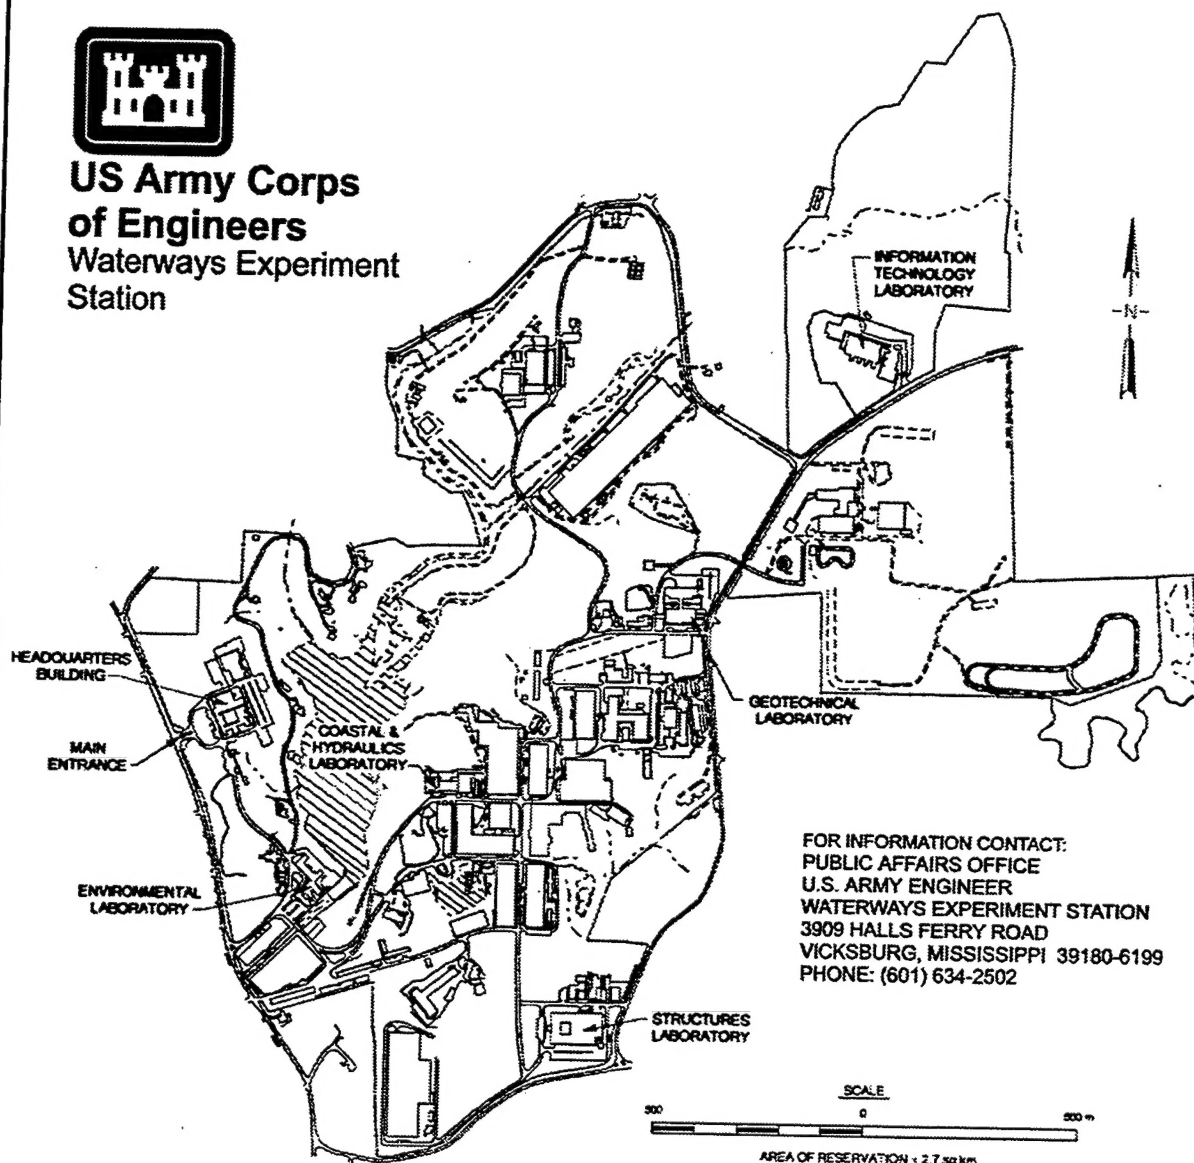
Preliminary report

Approved for public release; distribution is unlimited

DTIC QUALITY INSPECTED 8



**US Army Corps
of Engineers**
Waterways Experiment
Station



Waterways Experiment Station Cataloging-in-Publication Data

West, John E.

U.S. Army Engineer Waterways Experiment Station major shared resource center benchmark results and analysis : performance level 1 / by John E. West, Alex R. Carrillo ; prepared for U.S. Army Corps of Engineers.

41 p. : ill. ; 28 cm. -- (Technical report ; ITL-97-3)

Includes bibliographic references.

1. High performance computing -- United States -- Army. 2. Benchmarking (Management) -- United States -- Army. 3. System analysis -- United States -- Army. 4. Operations research -- United States -- Army. I. Carrillo, Alejandro R. II. United States. Army. Corps of Engineers. III. U.S. Army Engineer Waterways Experiment Station. IV. Information Technology Laboratory (U.S. Army Engineer Waterways Experiment Station) V. Title. VI. Series: Technical report (U.S. Army Engineer Waterways Experiment Station) ; ITL-97-3.

TA7 W34 no.ITL-97-3

Contents

Preface	iv
Executive Summary	v
1 Background	1
1.1 The Nichols Research Corporation Proposal	1
2 SGI Power Challenge Array	3
2.1 System Configuration	3
2.2 Witnessed Benchmark	3
2.3 Results	3
2.4 Recommendations	5
3 CRAY T3E	9
3.1 System Configuration	9
3.2 Witnessed Benchmark	10
3.3 Results	10
3.3.1 BM05 and BM06 Changes	11
3.3.2 BM20 and BM21 Changes	13
3.3.3 BM29	13
3.4 Recommendations	13
4 Summary and Recommendations	14
A Supporting Data: Power Challenge Array	A1
A.1 PCA1 Configuration	A2
A.2 PCA2 Configuration	A5
A.3 PCA1 Benchmark Iteration Comparison	A8
A.4 PCA2 Benchmark Iteration Comparison	A11
A.5 PCA1 Benchmark Staging Order	A13
A.6 PCA2 Benchmark Staging Order	A14
B Supporting Data: CRAY T3E	B1
B.1 Proposed CRAY T3D Configuration	B2
B.2 Installed CRAY T3E Configuration	B4

List of Figures

2.1	PCA1 proposed and witnessed benchmark staging and duration.	6
2.2	PCA2 proposed and witnessed benchmark staging and duration.	7
3.1	CRAY T3D (proposed) and CRAY T3E (witnessed) benchmark staging and duration.	12

List of Tables

1.1	NRC proposal benchmark times.	2
2.1	Benchmark iterations for each chassis.	4
2.2	Power Challenge Array witnessed benchmark times.	4
3.1	Benchmark iterations for the CRAY T3E and measured time.	10

Preface

This report¹ presents results and analysis of the Performance Level 1 target configuration benchmark for the DoD High Performance Computing (HPC) Major Shared Resource Center (MSRC) at the U.S. Army Engineer Waterways Experiment Station (WES). The MSRC at WES is operated as part of the DoD HPC Modernization Program of the Director, Defense Research and Engineering. Performance Level 1 refers to the first major phase of enhancements being made to the capabilities and capacities of the MSRC by Nichols Research Corporation, the integration contractor for the MSRC.

This work was performed by John E. West and Alex R. Carrillo, DoD High Performance Computing Center, Information Technology Laboratory (ITL), WES, Vicksburg, MS. The work was under the direction of Dr. N. Radhakrishnan, Director, ITL.

During preparation of this report, Dr. Robert W. Whalin was Director of WES. COL Bruce K. Howard, EN, was Commander.

¹The contents of this report are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval of the use of such commercial products.

Preliminary Investigation Executive Summary

Upon installation of the Performance Level 1 target configuration, the government's integration contract requires that the contractor demonstrate that the performance levels achieved by the installed configuration match those cited in the offeror's proposal for the Major Shared Resource Center (MSRC) target configuration High Performance Computing (HPC) resources. This demonstration, referred to herein as the "witnessed benchmark," has been completed, and preliminary results indicate several areas of concern.

The proposed configuration consisted of an SGI Power Challenge Array, a CRAY T3D, and a CRAY C90. The installed configuration consists of an SGI Power Challenge Array, a CRAY T3E, and a CRAY C90. The impact of the substitution of the CRAY T3E on code changes has been assessed and found to be minimal. Furthermore, the Measured Benchmark Time of the installed configuration, 6317 seconds, is less than the time for the proposed configuration, 6338 seconds.

However, analysis of the supplied benchmark data indicates that several concerns must be addressed before the benchmark can be recommended as successfully completed:

Power Challenge Array

- The I/O benchmark for the SGI Power Challenge Array must be rerun to demonstrate that the system can achieve the 100 MByte/sec transfer rate required in Section C.5.1.1.2.4 of the RFP.
- The government must have the opportunity to witness the execution of a single iteration of BM27 to resolve questions over the origin of the executable for this benchmark.
- A written explanation for the longer than proposed run time of the witnessed Power Challenge Array benchmark must be provided for further assessment.

CRAY T3E

- Although the changes made to the benchmarks for the CRAY T3E are few and are not found to introduce significant new complexity, they seem primarily directed to performance enhancements and not porting issues. The contractor must provide an explanation and justification for these changes.
- Evidence of I/O benchmark execution on the CRAY T3E was not found; such evidence is required and must be supplied by the contractor before final acceptance of this system.
- There is a discrepancy between the session log file which recorded the full details of the entire witnessed benchmark process and the supporting data and individual logs for BM29. The session log indicates a failed compilation due to a streams variable, while the individual log for BM29 does not indicate a problem. The contractor must provide an explanation of this discrepancy.
- The government must have the opportunity to witness the execution of a single iteration of BM29 to resolve questions over the origin of the executable for this benchmark.

Chapter 1

Background

The benchmarking process for the Department of Defense (DoD) Major Shared Resource Center (MSRC) procurement was designed to measure the ability of the offeror's proposed technical solution to satisfy the government's computational workload. For purposes of this contract, the workload is represented by a suite of codes assembled by the government from candidate applications submitted in each of the DoD Computational Technology Areas (CTAs). Offerors proposing a solution for a particular MSRC were then required to submit benchmark results for the set of codes representing the CTAs for which that MSRC was responsible. It was the offeror's responsibility to select a range of High Performance Computing (HPC) architectures and stage the benchmark applications on those machines to demonstrate that the technical configuration provided "best value" to the government. No guidance was given to the offeror concerning application staging or the types of architectures desired in the final solution - any solution which met the specific requirements for service in the Request for Proposal (RFP) was acceptable. To ensure that the government purchased a definable level of hardware performance, each offeror was required to abide by the processing and code change rules set forth in Attachment VI of the RFP. These rules contain limitations on the types of code changes which are allowable and define the information which must be provided to support the reported performance. The most important information which was to be provided in this regard was a listing of the codes changes with explanation and justification. The performance of the proposed target configuration is represented by the Measured Benchmark Time (MBT), computed as the longest of the elapsed (wall clock) times required for each individual measured HPC system to complete the benchmark iterations assigned to it. The magnitude of the MBT was assessed (by the Source Selection Evaluation Board) relative to the number and complexity of the code changes necessary to produce it, with special emphasis on performance "achievability" by "typical" DoD researchers. This qualitative assessment was then used to rate the offeror's technical solution.

1.1 The Nichols Research Corporation Proposal

After completion of the Source Selection Evaluation, Nichols Research Corporation (NRC) was awarded the contract for the MSRC at the U. S. Army Corps of Engineers Waterways Experiment Station (WES) MSRC. The proposal included the existing CRAY C90 (for which the offerors used government performance numbers and is thus not considered in this report), a new 32-node two-chassis Silicon Graphics, Inc. (SGI) Power Challenge Array (PCA), and a new 256-node Cray Research CRAY T3D with a CRAY Y-MP front-end. Table 1.1 shows the dates and system times of

System	Date	Benchmark Time (s)
SGI PCA	12/04/95	5004
CRAY T3D	05/04/95	6338
Measured Benchmark Time		6338

Table 1.1: NRC proposal benchmark times.

the final submitted benchmarks for each new system, along with the reported Measured Benchmark Time. The benchmark process executed by the offeror used the same model CRAY T3D and PCA proposed in Tab 4 of the proposal in addition to the CRAY C90, which was paper benchmarked as proscribed in the RFP, Attachment VI, paragraph 17. Pre-award benchmarks for the CRAY T3D were run at Cray Research, Inc., in Eagan, MN. PCA benchmarks were run at SGI in Mountain View, CA.

The RFP called for the installed MSRC systems to be re-benchmarked to ensure that the level of performance in the delivered configuration would match the performance of the proposed configuration. Nominally, this process consists of executing the benchmarks as submitted with the final pre-award proposal. However, this process is somewhat complicated in the WES case because the installed PL1 configuration differs from the proposed solution. Through the Engineering Change Proposal mechanism in the contract, the contractor substituted new computational technology, a CRAY T3E, for the proposed CRAY T3D (the remaining hardware is as proposed). The implications of this change to the benchmarking evaluation process and its impact on government acceptance of the installed PL1 configuration will be discussed in the CRAY T3E section of this report. First, the more straightforward benchmark and analysis for the installed PCA are discussed.

Chapter 2

SGI Power Challenge Array

2.1 System Configuration

The proposed and installed PCA systems are the same; the system is composed of two HiPPI-connected¹ chassis each containing sixteen 90 MHz R8000 MIPS processors with a 4 MByte secondary instruction/data cache, a 16 KByte data cache, a 16 KByte instruction cache, and 8192 MBytes of 8-way interleaved memory. The operating system is IRIX 6.2. Detailed information on the components in each chassis can be found in the system configuration tables in Appendix A.

2.2 Witnessed Benchmark

For minimal impact on MSRC operations, each chassis of the PCA was benchmarked separately. This posed no technical problems as the offeror did not spread any single iteration of a benchmark over processors in both chassis, effectively allowing the systems to be separated and each chassis benchmarked separately. The benchmark for the first chassis, PCA1, was performed on January 29, 1997 in the Joint Computing Facility of Building 8000 at WES beginning at 13:37 and ending at 15:16. The benchmark for the second chassis, PCA2, was performed on January 31, 1997 at the same location beginning at 13:35 and ending at 14:57. In each case Mitch Baker of NRC executed the benchmark which was witnessed by V. Sotler, J. West, and A. Carrillo, each government employees of WES.

At the beginning of each session, the NRC representative demonstrated the configuration of the system and that all directories and environment variables contained only expected data. The benchmark script was then started, beginning execution of the benchmark suite. During execution government personnel monitored the system to ensure that all processing rules were followed. Following execution the NRC representative generated a 4mm DAT tape containing all relevant information. These tapes were then transferred to the government for analysis.

2.3 Results

The data on the DAT tapes (two tapes, one for each chassis) were then examined to ensure that the rules and guidelines set forth in Attachment VI for benchmark processing and reporting were

¹HiPPI is the acronym for High Performance Parallel Interface, a network technology which supports high bandwidth connections between closely-coupled computers.

Benchmark	Required Iterations	PCA1	PCA2
01	13‡	5	5
02	12‡	0	6
03	20	20	0
23	9	4	5
24	11	11	0
25	44	44	0
27	35	0	35
30	24	24	0
31	5	3	2

‡ Balance of required iterations run on the C90

Table 2.1: Benchmark iterations for each chassis.

Chassis	Date	Benchmark Time (s)
PCA1	01/29/97	5581
PCA2	01/31/97	4902
System Benchmark Time		5581

Table 2.2: Power Challenge Array witnessed benchmark times.

followed. The government determined that in each case NRC ran the same benchmarks, in the same order, in the witnessed configuration as in the final pre-award proposal. Table 2.1 shows what these benchmarks were and that the proper number of iterations (as set forth in Attachment VI) of each were performed.

NRC did not deviate from the code changes submitted in the final proposal, and correctness criteria for each benchmark were met. In general, NRC performed the witnessed benchmark in accordance with all rules and met all requirements. In particular, no hardware or software (configuration) changes were made during the benchmark suite execution on either chassis, and neither machine was rebooted, reconfigured, or reinitialized at any time during the measured benchmark process. Table 2.2 shows the benchmark times for each chassis, and the reported system time.

There are some inconsistencies which need to be discussed. First, the pre-award proposal reported a time of 5004 seconds elapsed time for the PCA system (longest time for both chassis). The witnessed benchmark execution time was 5581 seconds. The tables in Appendix A show the percentage difference between the proposed and witnessed execution times for each iteration of each benchmark. Figures 2.1 and 2.2 show the start and stop times for each iteration for both the pre-award proposal and the witnessed benchmark runs (in these figures the y-axis quantity is derived by concatenating the benchmark number with individual iteration numbers; thus, the times for BM31 iteration 5 appear on the y-axis at 31.5). Analysis of these graphs and tables shows that the majority of the benchmark iterations run longer in the witnessed configuration than the proposed configuration, some by as much as 14% (BM01). However, these are wall clock times, and as such are tremendously sensitive. It is not uncommon to see variations of several percent in wall clock times between back-to-back executions of the same code; in this case the situation

is further complicated by variations in software versions (the contractor is required to provide the latest versions of compilers and operating system software upon installation; in the time between the final pre-award proposal and installation, several software products transitioned from beta to production versions) and the vagaries of machine configuration. It is thus difficult to establish a direct causal relationship between a particular circumstance and the increased benchmark time for this system. In this regard, however, it is useful to recall that the RFP states that the offeror is rated only on the Measured Benchmark Time, reported as the longest time among all of the measured HPC systems. The proposal indicates that this time is 6338 seconds (achieved by the CRAY T3D). The reported PCA time of 5581 seconds, while longer than the time of 5004 seconds reported in the proposal for the same system, is still significantly less than the Measured Benchmark Time for the MSRC.

Second, an anomaly was found in the execution of BM27 on PCA2. The session log shows a duration of 2 seconds for the edit/compile/link (ECL) stage during the witnessed benchmark, while the time reported in the proposal for the same events is 256 seconds. Further investigation of the ECL log found in SRC/WORK.N.20_x.chas.02/BM27 revealed that, when an attempt was made to compile the application, the file BM27.exe was already present. A listing of that file from the witnessed benchmark tape shows that the program was in fact compiled the morning of the benchmark as follows:

393632 Jan 31 10:47 BM27.exe*

As the benchmark session did not start until 13:30 on that day, it is clear that this program was compiled before the witnessed session. All other benchmark ECL stages were executed properly, and close examination revealed that all other required files for this benchmark (output, etc.) were created in the correct time frame. It is thought that this executable was left as an oversight from a test run conducted the morning before the witnessed benchmark. The benchmark staging order table for PCA2 in Appendix A shows that BM27 is the last benchmark staged on this system. Addition of the compile time for this benchmark (approximately 254 additional seconds from information given in benchmark iteration comparison table in Appendix A) would result in delaying the completion time for this benchmark by approximately four minutes for a final time on this chassis of approximately 5156 seconds – still less than the time for PCA1.

Finally, the I/O benchmark reported in the proposal achieved a transfer rate of 112.28 MByte/sec using two iterations (one per chassis) and 3200 MByte files. The witnessed transfer rate on the installed system was 87.67 MByte/sec. Section C.5.1.1.2.4, Disk Subsystems, of the RFP requires that the “aggregate data transfer rate across all disk subsystems on each HPC computer system...shall be a minimum of 100 MByte/sec.” In this instance the I/O benchmark does not demonstrate satisfaction of this minimum mandatory requirement, a situation which must be addressed by the contractor.

2.4 Recommendations

There are several technical concerns which must be addressed before we recommend the SGI Power Challenge Array component of the MSRC target configuration benchmark as successfully completed. First, subject to further discussion, the I/O benchmark should be rerun to demonstrate that the system can achieve a 100 MByte/sec transfer rate as required in Section C.5.1.1.2.4 of the RFP. Second, the government must evaluate the contractor's explanation for the longer than

PCA1 Benchmark Staging




Figure 2.1: PCA1 proposed and witnessed benchmark staging and duration.

PCA2 Benchmark Staging




Figure 2.2: PCA2 proposed and witnessed benchmark staging and duration.

proposed run time of the witnessed benchmark on this system. Finally, the government must have the opportunity to witness the execution of a single iteration of BM27 to resolve questions over the origin of the executable for this benchmark.

Chapter 3

CRAY T3E

3.1 System Configuration

As discussed in the introduction, the offeror did not propose or provide benchmark data for a CRAY T3E in the original proposal. The proposed system was a 256-node CRAY T3D with a CRAY Y-MP front-end. This system was substituted after award with an 80-node CRAY T3E (in fact, the installed system is a 256-node CRAY T3E, but only 80 of those nodes were subject to PL1 benchmarking). A detailed listing of hardware components for this system can be found in Appendix B, along with a listing from the proposal specifying the configuration of the proposed CRAY T3D.

The 80-node CRAY T3E was determined to be “equivalent” to a 256-node CRAY T3D by the contractor based on efficiency factors for the processors in each machine. The CRAY T3D has a peak processing capability of 0.15 GFLOPs per processor; through the benchmarking process NRC determined that the peak performance obtained on the government’s workload (as represented by the benchmark suite) was 0.0225 GLOPs per processor, an efficiency of $0.0225/0.15 = 0.15$. The vendor reports a peak performance of each CRAY T3E processor as 0.5 GLOPs, or a factor of three beyond the CRAY T3D processors. NRC then assumed the same efficiency factor of 0.15 for the government’s workload on the CRAY T3E, and found an estimated sustained performance on the government’s workload of 0.075 GLOPs per processor. The difference in projected sustained performance between these machines is thus approximately a factor of three, yielding a rough equivalence of 256 CRAY T3D nodes to 80 CRAY T3E nodes.

The substitution is found to be technically rational given the stated assumptions, though it does create potential complications for the benchmarking of the installed PL1 configuration. The largest potential complication arises from the fact that the proposed MSRC configuration performance was evaluated using the code changes and associated performance for the CRAY T3D system. If code changes on the new CRAY T3E deviated substantially in complexity from those submitted on the CRAY T3D, the issue of relative valuation versus the original would have to be addressed. As shown below, however, there were no substantial code changes made to the proposed CRAY T3D benchmarks in preparation for execution on the CRAY T3E.

One technical issue with the CRAY T3E hardware did arise during the benchmarking process. There is a known problem with the “streams” software interface. Under certain conditions this problem can cause processors to access old copies of data recently updated, leading to computation errors and, in some cases, to a deadlock of the entire machine. The current vendor solution to this problem is to disable the streams interface entirely, which subsequently reduces data access speed

Benchmark	Required Iterations	CRAY T3E
05	10	10
06	3	3
20	22	22
21	5	5
29	13	13
System Benchmark Time		6317

Table 3.1: Benchmark iterations for the CRAY T3E and measured time.

and can adversely affect performance. As a result, NRC was required to run the benchmark suite with the machine in the "streams disabled" mode.

3.2 Witnessed Benchmark

The benchmark for the CRAY T3E system was performed on February 28, 1997 in the Joint Computing Facility of Building 8000 at WES beginning at 10:21 and ending at 13:25. Dave Anderson of Cray Research executed the benchmark for NRC, which was witnessed by V. Sotler, a government employee of WES.

At the beginning of the session the NRC representative demonstrated the configuration of the system and that all directories, environment variables, and file systems contained only expected data. The benchmark script was then started, loading all iterations and compile operations into the system queue. The queue was then started, beginning execution of the benchmark suite. During execution, government personnel monitored the system to ensure that all processing rules were followed. Following execution the NRC representative generated a tar file which was subsequently archived on the MSRC's mass storage system. A 4mm DAT tape was made of this information for archival purposes.

3.3 Results

These data were then examined to ensure that the rules and guidelines set forth in Attachment VI for benchmark processing and reporting were followed. NRC ran the same benchmarks in the witnessed configuration as in the final proposal. Table 3.1 shows what these benchmarks were and that the proper number of iterations (as set forth in Attachment VI) of each were performed.

During the witnessed benchmark, no hardware or software (configuration) changes were made during the benchmark suite execution, and the machine was not rebooted, reconfigured, or reinitialized at any time during the measured benchmark process. Furthermore, all benchmark results satisfied their correctness criteria. However, no evidence was found in the submitted benchmark data that the I/O benchmark was executed; this data is required. Table 3.1 shows that the reported system time, 6317 seconds, is less than the proposal time of 6338 seconds. Figure 3.1 shows the start and stop times for each iteration for both the proposal and the witnessed benchmark runs.

During analysis of the benchmark data, it was found that NRC did deviate somewhat from the code changes submitted in the final proposal. This was anticipated given substitution of a CRAY T3E for the CRAY T3D. These changes, detailed in sections 3.3.1 through 3.3.3, are not viewed as

particularly numerous or complex, and do not add to the overall complexity of modifications to the codes or their achievability by a "typical" researcher. However, some of the changes, particularly to BM20, BM21, and BM29 appear to be solely for the purpose of added optimization and not directed at porting the code to a new architecture. The contractor must supply a justification and explanation for these changes and assess their impact on code performance.

There is one final area of concern. In analysis of the supplied benchmark session log found in the file `cewes_t3e.2.script`, an apparent error in compilation of BM29 was discovered. Specifically:

```
Applying patch: patch_t3d/bench.f.t3d...
Applying patch: patch_t3d/evolve.f.t3d...
Applying patch: patch_t3d/include.h.t3d...
Applying patch: patch_t3d/io.f.t3d...
Applying patch: patch_t3d/main.f.t3d...
Applying patch: patch_t3d/subs.f.t3d...
    f90 -c -dp -O0 bench.f
    f90 -c -dp -O1,unroll2,nopattern evolve.f
    f90 -c -dp -O0 io.f
    f90 -c -dp -O1,unroll2,nopattern main.f
    f90 -c -dp -O1,unroll2,nopattern subs.f
    f90 -X32 -Wl-Dstreams=on -o ../bin_t3d/bm29.32 bench.o evolve.o
                                io.o main.o subs.o

cld-319 cld: ERROR
    The value 'on' for the directive 'streams' is not valid.
cld-117 cld: FATAL
    Errors occurred processing the input files.
make: "f90 -X32 -Wl-Dstreams=on -o ../bin_t3d/bm29.32 bench.o evolve.o
    io.o main.o subs.o": Error code 1
cmd-2436 make: Stop.
```

This error causes concern for three reasons. First, the indication in the script of a failed compilation raises questions over the origin of the executable used to run the benchmark. Second, the directive which caused the failed compilation seems designed to activate streams, which must be off for the witnessed benchmark of this system. Finally, it represents a discrepancy between the script file which (presumably) recorded all output of the benchmark session (`cewes_t3e.2.script`), and the remainder of the logs and supporting data provided for BM29. None of this supporting data indicates a failed compilation or the use of the failed loader directive. The contractor must provide an explanation for this discrepancy.

3.3.1 BM05 and BM06 Changes

The only modifications to this benchmark from those submitted for the CRAY T3D were to the routine `Admin_Luns.f`. This routine is used to coordinate the assignment of logical unit identifiers for file processing, and has no effect on the order of computation or on run times. The changes converted the original code segment from a single routine with multiple entry points to multiple routines (each with a single entry point). Also, an external declaration was removed for the function `PvmMin`.

T3E Benchmark Staging




Figure 3.1: CRAY T3D (proposed) and CRAY T3E (witnessed) benchmark staging and duration.

3.3.2 BM20 and BM21 Changes

The only modification to this benchmark for the CRAY T3E was separation of the routine MHMTB0 into its own file (MHMTB0.F).

3.3.3 BM29

The primary modifications to this benchmark from those proposed for the CRAY T3D involved changes due to memory size. Various memory variables in bench.f (mem_pe, 1 change) and main.f (various statements, 5 changes) were doubled. Also, the inner loop (dimension 14) of the subroutine equilibrate (found in the file evolve.f) was unrolled 14 times.

3.4 Recommendations

There are several technical concerns which must be addressed before we recommend the CRAY T3E component of the MSRC target configuration benchmark as successfully completed. The Measured Benchmark Time for the CRAY T3E is under the proposal MSRC Benchmark Time, and the changes made to the benchmark suite beyond those proposed for the CRAY T3D are few and relatively straightforward. In general, the modifications do not appear to introduce any significant changes in the way the code was executed (memory models, etc.) or in the complexity of the resulting codes. However, the government must have the opportunity to review the contractor's explanation and justification for these changes, as they seem to be focused primarily on performance rather than porting. Also, evidence of I/O benchmark execution on this system was not found; such evidence is required and must be supplied by the contractor before final acceptance of this system. Finally, an explanation of the discrepancy between the session log file and the supporting data for BM29 must be provided, and the government must witness the execution of a single iteration of this benchmark to resolve questions over the origin of its executable.

Chapter 4

Summary and Recommendations

The witnessed benchmark of the contractor-installed WES PL1 HPC target configuration has been completed. The contractor proposed an HPC target configuration consisting of a 32-processor SGI Power Challenge Array, a 256-processor CRAY T3D, and a 16-processor CRAY C90. Through the Engineering Change Proposal process, the CRAY T3D component of the configuration was replaced with an 80-processor CRAY T3E. The impact of the substitution of the CRAY T3E on code changes has been assessed and appears to be minimal. Furthermore, the Measured Benchmark Time of the installed configuration, 6317 seconds, is less than the time for the proposed configuration, 6338 seconds.

Analysis of the supplied benchmark data indicates that several concerns must be addressed before the benchmark can be recommended as successfully completed. A detailed discussion of these concerns and the issues surrounding them may be found in the preceding chapters.

Power Challenge Array

- The I/O benchmark for SGI Power Challenge Array must be rerun to demonstrate that the system can achieve the 100 MByte/sec transfer rate required in Section C.5.1.1.2.4 of the RFP.
- The government must have the opportunity to witness the execution of a single iteration of BM27 to resolve questions over the origin of the executable for this benchmark.
- A written explanation for the longer than proposed run time of the witnessed Power Challenge Array Benchmark must be provided for further assessment.

CRAY T3E

- Although the changes made to the benchmarks for the CRAY T3E are few and are not found to introduce significant new complexity, they seem primarily directed to performance enhancements and not porting issues. The contractor must provide an explanation and justification for these changes.
- Evidence of I/O benchmark execution on the CRAY T3E was not found; such evidence is required and must be supplied by the contractor before final acceptance of this system.

- There is a discrepancy between the session log file which recorded the full details of the entire benchmark process and the supporting data and individual logs for BM29. The session log indicates a failed compilation due to a streams variable, while the individual for BM29 does not indicate a problem. The contractor must provide an explanation of this discrepancy.
- The government must have the opportunity to witness the execution of a single iteration of BM29 to resolve questions over the origin of the executable for this benchmark.

Appendix A

Supporting Data: Power Challenge Array

A.1 PCA1 Configuration

The following is selected output from the `sysconf` command executed on PCA1.

```
VENDOR                Silicon Graphics, Inc.
OS_PROVIDER            Silicon Graphics, Inc.
OS_NAME                IRIX64
HW_NAME                IP21
NUM_PROCESSORS          16
HOSTID                 86a40d15
OSREL_MAJ               6
OSREL_MIN               2
OSREL_PATCH             0
PROCESSORS              R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0,
                        R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0,
                        R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0,
                        R8000 3.0, R8000 3.0
AVAIL_PROCESSORS        16
SYSNAME                 IRIX64
HOSTNAME                pca1
RELEASE                 6.2
VERSION                 03131015
ACHINE                  IP21
ARCHITECTURE             mips
HW_SERIAL                3442723
HW_PROVIDER              sgi
SRPC_DOMAIN              Not supported
INITTAB_NAME             /etc/inittab
```

The following information is extracted from the output of the `hinv` command executed on PCA1.

```
16 90 MHZ IP21 Processors
CPU: MIPS R8000 Processor Chip Revision: 3.0
FPU: MIPS R8010 Floating Point Chip Revision: 0.2
Secondary unified instruction/data cache size: 4 Mbytes
Data cache size: 16 Kbytes
Instruction cache size: 16 Kbytes
ain memory size: 8192 Mbytes, 8-way interleaved
I/O board, Ebus slot 11: IO4 revision 1
I/O board, Ebus slot 15: IO4 revision 1
Integral EPC serial ports: 4
Integral Ethernet controller: et0, Ebus slot 15
XPI FDDI controller: xpi0, slot 11, adapter 13, firmware version 9603091500, DAS
```

XPI FDDI controller: xpi1, slot 11, adapter 13, firmware version 9603091500, DAS
EPC external interrupts

Integral SCSI controller 111: Version WD33C95A, differential, revision 0
Disk drive: unit 4 on SCSI controller 111
Disk drive: unit 2 on SCSI controller 111
Disk drive: unit 1 on SCSI controller 111

Integral SCSI controller 110: Version WD33C95A, differential, revision 0
Disk drive: unit 4 on SCSI controller 110
Disk drive: unit 3 on SCSI controller 110
Disk drive: unit 2 on SCSI controller 110
Disk drive: unit 1 on SCSI controller 110

Integral SCSI controller 1: Version WD33C95A, differential, revision 0
Disk drive: unit 6 on SCSI controller 1
Disk drive: unit 5 on SCSI controller 1
Disk drive: unit 4 on SCSI controller 1
Disk drive: unit 3 on SCSI controller 1
Disk drive: unit 2 on SCSI controller 1
Disk drive: unit 1 on SCSI controller 1

Integral SCSI controller 0: Version WD33C95A, single ended, revision 0
CDROM: unit 5 on SCSI controller 0
Tape drive: unit 1 on SCSI controller 0: DAT

Integral SCSI controller 7: Version SCIP/WD33C95A, differential
Disk drive: unit 5 on SCSI controller 7
Disk drive: unit 3 on SCSI controller 7
Disk drive: unit 2 on SCSI controller 7
Disk drive: unit 1 on SCSI controller 7

Integral SCSI controller 6: Version SCIP/WD33C95A, differential
Disk drive: unit 3 on SCSI controller 6
Disk drive: unit 2 on SCSI controller 6
Disk drive: unit 1 on SCSI controller 6

Integral SCSI controller 5: Version SCIP/WD33C95A, differential
Disk drive: unit 3 on SCSI controller 5
Disk drive: unit 2 on SCSI controller 5
Disk drive: unit 1 on SCSI controller 5

HIPPI adapter: hipp11, slot 11 adap 6, firmware version 3321952
HIPPI adapter: hipp10, slot 15 adap 5, firmware version 3321952
CC synchronization join counter

Integral EPC parallel port: Ebus slot 11
Integral EPC parallel port: Ebus slot 15
VME bus: adapter 0 mapped to adapter 61
VME bus: adapter 61

The following compilers were proposed and used to run the original benchmarks (selected output of the versions command):

I	ftn77_dev	10/13/95	Fortran 77, 6.1
I	ftn90_dev	10/13/95	Fortran 90, 6.2ALPHA
I	c_dev	10/13/95	C, 6.2ALPHA

The following compilers are installed and used to run the witnessed benchmarks (selected output of the `versions` command):

I	ftn77_dev	01/23/97	Fortran 77, 6.2
I	ftn90_dev	01/09/97	Fortran 90, 7.1 on irix 6.2
I	c_dev	01/23/97	C, 6.2

A.2 PCA2 Configuration

The following is selected output from the `sysconf` command executed on PCA2.

```
VENDOR                Silicon Graphics, Inc.
OS_PROVIDER            Silicon Graphics, Inc.
OS_NAME                IRIX64
HW_NAME                IP21
NUM_PROCESSORS         16
HOSTID                 86a40d16
OSREL_MAJ              6
OSREL_MIN              2
OSREL_PATCH            0
PROCESSORS             R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0,
                      R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0,
                      R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0, R8000 3.0,
                      R8000 3.0, R8000 3.0
AVAIL_PROCESSORS       16
PATH                   :/usr/sbin:/usr/bsd:/sbin:/usr/bin:/bin:/usr/bin/X11
CS_PATH                :/usr/sbin:/usr/bsd:/sbin:/usr/bin:/bin:/usr/bin/X11
_CS_PATH               :/usr/sbin:/usr/bsd:/sbin:/usr/bin:/bin:/usr/bin/X11
SYSNAME                IRIX64
HOSTNAME               pca2
RELEASE                6.2
VERSION                03131015
ACHINE                 IP21
ARCHITECTURE           mips
HW_SERIAL              3442724
HW_PROVIDER            sgi
SRPC_DOMAIN            Not supported
INITTAB_NAME           /etc/inittab
```

The following information is extracted from the output of the `hinv` command executed on PCA2.

```
16 90 MHZ IP21 Processors
CPU: MIPS R8000 Processor Chip Revision: 3.0
FPU: MIPS R8010 Floating Point Chip Revision: 0.2
Secondary unified instruction/data cache size: 4 Mbytes
Data cache size: 16 Kbytes
Instruction cache size: 16 Kbytes
ain memory size: 8192 Mbytes, 8-way interleaved
I/O board, Ebus slot 13: IO4 revision 1
I/O board, Ebus slot 15: IO4 revision 1
```

Integral EPC serial ports: 4
 Integral Ethernet controller: et0, Ebus slot 15
 XPI FDDI controller: xpi0, slot 13, adapter 13, firmware version 9603091500, DAS
 XPI FDDI controller: xpi1, slot 13, adapter 13, firmware version 9603091500, DAS
 EPC external interrupts
 Integral SCSI controller 131: Version WD33C95A, differential, revision 0
 Disk drive: unit 4 on SCSI controller 131
 Disk drive: unit 3 on SCSI controller 131
 Disk drive: unit 2 on SCSI controller 131
 Disk drive: unit 1 on SCSI controller 131
 Integral SCSI controller 130: Version WD33C95A, differential, revision 0
 Disk drive: unit 3 on SCSI controller 130
 Disk drive: unit 2 on SCSI controller 130
 Disk drive: unit 1 on SCSI controller 130
 Integral SCSI controller 1: Version WD33C95A, differential, revision 0
 Disk drive: unit 4 on SCSI controller 1
 Disk drive: unit 3 on SCSI controller 1
 Disk drive: unit 2 on SCSI controller 1
 Disk drive: unit 1 on SCSI controller 1
 Integral SCSI controller 0: Version WD33C95A, differential, revision 0
 Disk drive: unit 4 on SCSI controller 0
 Disk drive: unit 3 on SCSI controller 0
 Disk drive: unit 2 on SCSI controller 0
 Disk drive: unit 1 on SCSI controller 0
 Integral SCSI controller 4: Version SCIP/WD33C95A, differential
 Disk drive: unit 3 on SCSI controller 4
 Disk drive: unit 2 on SCSI controller 4
 Disk drive: unit 1 on SCSI controller 4
 Integral SCSI controller 3: Version SCIP/WD33C95A, differential
 Disk drive: unit 3 on SCSI controller 3
 Disk drive: unit 2 on SCSI controller 3
 Disk drive: unit 1 on SCSI controller 3
 Integral SCSI controller 2: Version SCIP/WD33C95A, differential
 Disk drive: unit 3 on SCSI controller 2
 Disk drive: unit 2 on SCSI controller 2
 Disk drive: unit 1 on SCSI controller 2
 HIPPI adapter: hipp11, slot 13 adap 6, firmware version 3321952
 HIPPI adapter: hipp10, slot 15 adap 6, firmware version 3321952
 CC synchronization join counter
 Integral EPC parallel port: Ebus slot 13
 Integral EPC parallel port: Ebus slot 15
 VME bus: adapter 0 mapped to adapter 61
 VME bus: adapter 61

The following compilers were proposed and used to run the original benchmarks (selected output of the versions command):

I	ftn77_dev	10/13/95	Fortran 77, 6.1
I	ftn90_dev	10/13/95	Fortran 90, 6.2ALPHA
I	c_dev	10/13/95	C, 6.2ALPHA

The following compilers are installed and used to run the witnessed benchmarks (selected output of the `versions` command):

I	ftn77_dev	01/23/97	Fortran 77, 6.2
I	ftn90_dev	01/09/97	Fortran 90, 7.1 on irix 6.2
I	c_dev	01/23/97	C, 6.2

A.3 PCA1 Benchmark Iteration Comparison

Proposal BM/Itr : t(s)	Witness BM/Itr : t(s)	Diff $100 \cdot (\text{wit} - \text{prop}) / \text{prop}$
BM01ECL001 : 365	BM01ECL001 : 366	+0.27
BM01RUN001 : 3345	BM01RUN001 : 3694	+10.43
BM01RUN002 : 3339	BM01RUN002 : 3785	+13.36
BM01RUN003 : 3233	BM01RUN003 : 3734	+15.50
BM01RUN004 : 3289	BM01RUN004 : 3741	+13.74
BM01RUN005 : 3235	BM01RUN005 : 3773	+16.63
BM03ECL001 : 82	BM03ECL001 : 90	+9.76
BM03RUN001 : 243	BM03RUN001 : 254	+4.53
BM03RUN002 : 242	BM03RUN002 : 245	+1.24
BM03RUN003 : 243	BM03RUN003 : 247	+1.65
BM03RUN004 : 241	BM03RUN004 : 248	+2.90
BM03RUN005 : 241	BM03RUN005 : 249	+3.32
BM03RUN006 : 244	BM03RUN006 : 258	+5.74
BM03RUN007 : 243	BM03RUN007 : 247	+1.65
BM03RUN008 : 244	BM03RUN008 : 248	+1.64
BM03RUN009 : 244	BM03RUN009 : 247	+1.23
BM03RUN010 : 244	BM03RUN010 : 245	+0.41
BM03RUN011 : 243	BM03RUN011 : 248	+2.06
BM03RUN012 : 242	BM03RUN012 : 245	+1.24
BM03RUN013 : 242	BM03RUN013 : 246	+1.65
BM03RUN014 : 242	BM03RUN014 : 250	+3.31
BM03RUN015 : 241	BM03RUN015 : 247	+2.49
BM03RUN016 : 242	BM03RUN016 : 245	+1.24
BM03RUN017 : 243	BM03RUN017 : 246	+1.23
BM03RUN018 : 241	BM03RUN018 : 247	+2.49
BM03RUN019 : 241	BM03RUN019 : 246	+2.07
BM03RUN020 : 242	BM03RUN020 : 247	+2.07
BM23ECL001 : 8	BM23ECL001 : 9	+12.50
BM23RUN001 : 1267	BM23RUN001 : 1387	+9.47
BM23RUN002 : 1257	BM23RUN002 : 1231	-2.07
BM23RUN003 : 1233	BM23RUN003 : 1269	+2.92
BM23RUN004 : 1246	BM23RUN004 : 1330	+6.74
BM24ECL001 : 6	BM24ECL001 : 9	+50.00
BM24RUN001 : 111	BM24RUN001 : 123	+10.81
BM24RUN002 : 112	BM24RUN002 : 129	+15.18
BM24RUN003 : 112	BM24RUN003 : 123	+9.82

BM24RUN004 : 112	BM24RUN004 : 126	+12.50
BM24RUN005 : 111	BM24RUN005 : 129	+16.22
BM24RUN006 : 113	BM24RUN006 : 129	+14.16
BM24RUN007 : 112	BM24RUN007 : 129	+15.18
BM24RUN008 : 111	BM24RUN008 : 128	+15.32
BM24RUN009 : 112	BM24RUN009 : 126	+12.50
BM24RUN010 : 112	BM24RUN010 : 128	+14.29
BM24RUN011 : 114	BM24RUN011 : 129	+13.16

BM25ECL001 : 3	BM25ECL001 : 5	+66.67
BM25RUN001 : 52	BM25RUN001 : 56	+7.69
BM25RUN002 : 53	BM25RUN002 : 55	+3.77
BM25RUN003 : 52	BM25RUN003 : 56	+7.69
BM25RUN004 : 51	BM25RUN004 : 57	+11.76
BM25RUN005 : 51	BM25RUN005 : 58	+13.73
BM25RUN006 : 53	BM25RUN006 : 58	+9.43
BM25RUN007 : 54	BM25RUN007 : 64	+18.52
BM25RUN008 : 53	BM25RUN008 : 57	+7.55
BM25RUN009 : 52	BM25RUN009 : 57	+9.62
BM25RUN010 : 51	BM25RUN010 : 51	+0.00
BM25RUN011 : 52	BM25RUN011 : 52	+0.00
BM25RUN012 : 51	BM25RUN012 : 51	+0.00
BM25RUN013 : 52	BM25RUN013 : 52	+0.00
BM25RUN014 : 51	BM25RUN014 : 52	+1.96
BM25RUN015 : 52	BM25RUN015 : 52	+0.00
BM25RUN016 : 51	BM25RUN016 : 51	+0.00
BM25RUN017 : 51	BM25RUN017 : 58	+13.73
BM25RUN018 : 52	BM25RUN018 : 52	+0.00
BM25RUN019 : 52	BM25RUN019 : 54	+3.85
BM25RUN020 : 52	BM25RUN020 : 53	+1.92
BM25RUN021 : 52	BM25RUN021 : 52	+0.00
BM25RUN022 : 52	BM25RUN022 : 52	+0.00
BM25RUN023 : 52	BM25RUN023 : 52	+0.00
BM25RUN024 : 52	BM25RUN024 : 53	+1.92
BM25RUN025 : 51	BM25RUN025 : 51	+0.00
BM25RUN026 : 52	BM25RUN026 : 52	+0.00
BM25RUN027 : 52	BM25RUN027 : 52	+0.00
BM25RUN028 : 52	BM25RUN028 : 52	+0.00
BM25RUN029 : 52	BM25RUN029 : 52	+0.00
BM25RUN030 : 51	BM25RUN030 : 52	+1.96
BM25RUN031 : 51	BM25RUN031 : 52	+1.96
BM25RUN032 : 51	BM25RUN032 : 53	+3.92
BM25RUN033 : 51	BM25RUN033 : 52	+1.96
BM25RUN034 : 51	BM25RUN034 : 52	+1.96
BM25RUN035 : 52	BM25RUN035 : 52	+0.00
BM25RUN036 : 51	BM25RUN036 : 52	+1.96

BM25RUN037 : 51	BM25RUN037 : 52	+1.96
BM25RUN038 : 51	BM25RUN038 : 51	+0.00
BM25RUN039 : 51	BM25RUN039 : 51	+0.00
BM25RUN040 : 51	BM25RUN040 : 51	+0.00
BM25RUN041 : 50	BM25RUN041 : 51	+2.00
BM25RUN042 : 50	BM25RUN042 : 51	+2.00
BM25RUN043 : 51	BM25RUN043 : 51	+0.00
BM25RUN044 : 51	BM25RUN044 : 52	+1.96
BM30ECL001 : 395	BM30ECL001 : 370	-6.33
BM30RUN001 : 213	BM30RUN001 : 228	+7.04
BM30RUN002 : 215	BM30RUN002 : 239	+11.16
BM30RUN003 : 214	BM30RUN003 : 233	+8.88
BM30RUN004 : 215	BM30RUN004 : 242	+12.56
BM30RUN005 : 215	BM30RUN005 : 235	+9.30
BM30RUN006 : 214	BM30RUN006 : 240	+12.15
BM30RUN007 : 216	BM30RUN007 : 230	+6.48
BM30RUN008 : 216	BM30RUN008 : 241	+11.57
BM30RUN009 : 214	BM30RUN009 : 227	+6.07
BM30RUN010 : 214	BM30RUN010 : 240	+12.15
BM30RUN011 : 215	BM30RUN011 : 240	+11.63
BM30RUN012 : 216	BM30RUN012 : 240	+11.11
BM30RUN013 : 217	BM30RUN013 : 231	+6.45
BM30RUN014 : 216	BM30RUN014 : 240	+11.11
BM30RUN015 : 213	BM30RUN015 : 241	+13.15
BM30RUN016 : 217	BM30RUN016 : 233	+7.37
BM30RUN017 : 215	BM30RUN017 : 222	+3.26
BM30RUN018 : 214	BM30RUN018 : 227	+6.07
BM30RUN019 : 214	BM30RUN019 : 224	+4.67
BM30RUN020 : 214	BM30RUN020 : 225	+5.14
BM30RUN021 : 214	BM30RUN021 : 223	+4.21
BM30RUN022 : 215	BM30RUN022 : 228	+6.05
BM30RUN023 : 215	BM30RUN023 : 225	+4.65
BM30RUN024 : 215	BM30RUN024 : 227	+5.58
BM31ECL001 : 328	BM31ECL001 : 374	+14.02
BM31RUN001 : 3393	BM31RUN001 : 3406	+0.38
BM31RUN002 : 3394	BM31RUN002 : 3468	+2.18
BM31RUN003 : 3382	BM31RUN003 : 3489	+3.16

A.4 PCA2 Benchmark Iteration Comparison

Proposal BM/Itr : t(s)	Witness BM/Itr : t(s)	Diff 100*(wit-prop)/prop
BM01ECL001 : 389	BM01ECL001 : 332	-14.65
BM01RUN001 : 3217	BM01RUN001 : 3529	+9.70
BM01RUN002 : 3260	BM01RUN002 : 3517	+7.88
BM01RUN003 : 3218	BM01RUN003 : 3528	+9.63
BM01RUN004 : 3219	BM01RUN004 : 3530	+9.66
BM01RUN005 : 3235	BM01RUN005 : 3533	+9.21
BM02ECL001 : 389	BM02ECL001 : 332	-14.65
BM02RUN001 : 2706	BM02RUN001 : 2768	+2.29
BM02RUN002 : 2696	BM02RUN002 : 2755	+2.19
BM02RUN003 : 2594	BM02RUN003 : 2756	+6.25
BM02RUN004 : 2708	BM02RUN004 : 2749	+1.51
BM02RUN005 : 2606	BM02RUN005 : 2761	+5.95
BM02RUN006 : 2616	BM02RUN006 : 2766	+5.73
BM23ECL001 : 9	BM23ECL001 : 8	-11.11
BM23RUN001 : 1257	BM23RUN001 : 1369	+8.91
BM23RUN002 : 1259	BM23RUN002 : 1497	+18.90
BM23RUN003 : 1260	BM23RUN003 : 1497	+18.81
BM23RUN004 : 1260	BM23RUN004 : 1306	+3.65
BM23RUN005 : 1248	BM23RUN005 : 1497	+19.95
BM27ECL001 : 256	BM27ECL001 : 2	-99.22
BM27RUN001 : 260	BM27RUN001 : 273	+5.00
BM27RUN002 : 261	BM27RUN002 : 275	+5.36
BM27RUN003 : 262	BM27RUN003 : 277	+5.73
BM27RUN004 : 264	BM27RUN004 : 276	+4.55
BM27RUN005 : 263	BM27RUN005 : 281	+6.84
BM27RUN006 : 274	BM27RUN006 : 277	+1.09
BM27RUN007 : 260	BM27RUN007 : 277	+6.54
BM27RUN008 : 262	BM27RUN008 : 269	+2.67
BM27RUN009 : 261	BM27RUN009 : 283	+8.43
BM27RUN010 : 260	BM27RUN010 : 269	+3.46
BM27RUN011 : 261	BM27RUN011 : 271	+3.83
BM27RUN012 : 263	BM27RUN012 : 270	+2.66
BM27RUN013 : 261	BM27RUN013 : 275	+5.36
BM27RUN014 : 262	BM27RUN014 : 272	+3.82
BM27RUN015 : 260	BM27RUN015 : 274	+5.38
BM27RUN016 : 262	BM27RUN016 : 275	+4.96

BM27RUN017 : 260	BM27RUN017 : 270	+3.85
BM27RUN018 : 262	BM27RUN018 : 273	+4.20
BM27RUN019 : 261	BM27RUN019 : 274	+4.98
BM27RUN020 : 262	BM27RUN020 : 274	+4.58
BM27RUN021 : 280	BM27RUN021 : 280	+0.00
BM27RUN022 : 265	BM27RUN022 : 280	+5.66
BM27RUN023 : 260	BM27RUN023 : 274	+5.38
BM27RUN024 : 261	BM27RUN024 : 272	+4.21
BM27RUN025 : 261	BM27RUN025 : 272	+4.21
BM27RUN026 : 261	BM27RUN026 : 270	+3.45
BM27RUN027 : 261	BM27RUN027 : 275	+5.36
BM27RUN028 : 260	BM27RUN028 : 270	+3.85
BM27RUN029 : 261	BM27RUN029 : 270	+3.45
BM27RUN030 : 261	BM27RUN030 : 271	+3.83
BM27RUN031 : 260	BM27RUN031 : 269	+3.46
BM27RUN032 : 269	BM27RUN032 : 276	+2.60
BM27RUN033 : 261	BM27RUN033 : 279	+6.90
BM27RUN034 : 261	BM27RUN034 : 269	+3.07
BM27RUN035 : 260	BM27RUN035 : 270	+3.85
BM31ECL001 : 319	BM31ECL001 : 310	-2.82
BM31RUN001 : 3403	BM31RUN001 : 3347	-1.65
BM31RUN002 : 3331	BM31RUN002 : 3336	+0.15
BM31RUN002 : 3336	BM31RUN002 : 3336	+0.00

A.5 PCA1 Benchmark Staging Order

The following is a selection of text from the script which controls the execution order of each benchmark iteration to be run on this chassis. Benchmarks in the witnessed configuration were staged in the same order as those in the proposal.

```
alljobs = BM31ECL001 BM01ECL001 BM31RUN001 BM31RUN002 BM31RUN003
BM01RUN003 BM01RUN004 BM01RUN001 BM01RUN002 BM01RUN005 BM23ECL001
BM23RUN001 BM23RUN002 BM23RUN003 BM23RUN004 BM03ECL001 BM03RUN003
BM03RUN001 BM03RUN002 BM03RUN005 BM03RUN006 BM03RUN007 BM03RUN008
BM03RUN009 BM03RUN010 BM03RUN011 BM03RUN012 BM03RUN013 BM03RUN016
BM03RUN017 BM03RUN018 BM03RUN019 BM03RUN020 BM03RUN004 BM03RUN014
BM03RUN015 BM30ECL001 BM30RUN001 BM30RUN002 BM30RUN003 BM30RUN004
BM30RUN005 BM30RUN006 BM30RUN007 BM30RUN008 BM30RUN009 BM30RUN010
BM30RUN011 BM30RUN012 BM30RUN013 BM30RUN014 BM30RUN015 BM30RUN016
BM30RUN017 BM30RUN018 BM30RUN019 BM30RUN020 BM30RUN021 BM30RUN022
BM30RUN023 BM30RUN024 BM24ECL001 BM24RUN001 BM24RUN002 BM24RUN003
BM24RUN004 BM24RUN005 BM24RUN006 BM24RUN007 BM24RUN008 BM24RUN009
BM24RUN010 BM24RUN011 BM25ECL001 BM25RUN001 BM25RUN002 BM25RUN003
BM25RUN004 BM25RUN005 BM25RUN006 BM25RUN007 BM25RUN008 BM25RUN009
BM25RUN010 BM25RUN011 BM25RUN012 BM25RUN013 BM25RUN014 BM25RUN015
BM25RUN016 BM25RUN017 BM25RUN018 BM25RUN019 BM25RUN020 BM25RUN021
BM25RUN022 BM25RUN023 BM25RUN024 BM25RUN025 BM25RUN026 BM25RUN027
BM25RUN028 BM25RUN029 BM25RUN030 BM25RUN031 BM25RUN032 BM25RUN033
BM25RUN034 BM25RUN035 BM25RUN036 BM25RUN037 BM25RUN038 BM25RUN039
BM25RUN040 BM25RUN041 BM25RUN042 BM25RUN043 BM25RUN044
```

A.6 PCA2 Benchmark Staging Order

The following is a selection of text from the script which controls the execution order of each benchmark iteration to be run on this chassis. Benchmarks in the witnessed configuration were staged in the same order as those in the proposal.

```
alljobs = BM02ECL001 BM01ECL001 BM31ECL001 BM31RUN002 BM01RUN001
BM01RUN004 BM01RUN003 BM01RUN005 BM01RUN002 BM31RUN001 BM02RUN003
BM02RUN005 BM02RUN002 BM02RUN004 BM02RUN001 BM02RUN006 BM23ECL001
BM23RUN004 BM23RUN001 BM23RUN003 BM23RUN002 BM23RUN005 BM27ECL001
BM27RUN027 BM27RUN009 BM27RUN033 BM27RUN007 BM27RUN002 BM27RUN003
BM27RUN004 BM27RUN005 BM27RUN006 BM27RUN021 BM27RUN032 BM27RUN022
BM27RUN030 BM27RUN011 BM27RUN012 BM27RUN013 BM27RUN014 BM27RUN015
BM27RUN016 BM27RUN018 BM27RUN019 BM27RUN020 BM27RUN023 BM27RUN024
BM27RUN025 BM27RUN026 BM27RUN028 BM27RUN029 BM27RUN034 BM27RUN035
BM27RUN017 BM27RUN031 BM27RUN001 BM27RUN010 BM27RUN008
```

Appendix B

Supporting Data: CRAY T3E

B.1 Proposed CRAY T3D Configuration

The following is selected output from a log file provided with the CRAY T3D benchmark data.

Basic configuration was Y-MP 8E with 4 processors and 64 Mw with 12 channels of DD-60 disk drives.

The Y-MP is attached to a 256 node Cray T3D system with two high-speed connections.

...

This is a 4-CPU, 64 MW system

```
HARDWARE: SERIAL= SN1049   MFTYPE= CRAY-YMP
           MFSUBTYPE= YMPOXX   NCPU= 4   CPCYCLE= 6.0000 ns
           MEM= 67106560   NBANKS= 128   CHIPSZ= 262144
           AVL= YES   BDM= YES   EMA= YES   HPM= YES   BMM= NO
           SSD= 134217728   NVHISP= 2   IOS= MODEL_E

SOFTWARE:  RELEASE= 8.00   POSIX VERSION= 199009   SECURE SYS= OFF
           SYMEM= 10020864 WRDS   USRMEM= 57085696 WRDS
           OS_HZ= 60   CLK_TCK= 166666667
           JOB_CONTROL= YES   SAVED_IDS= YES   SCTRACE= ON
           UID_MAX= 60000   PID_MAX= 100000
           ARG_MAX= 49999   CHILD_MAX= 500   OPEN_MAX= 1024
           NMOUNT= 200   NUSERS= 200   NPTY= 255
           NDISK= 256   SDS= 0   NBUF= 5000
           PRIV_SU= ON   PRIV_TFM= OFF
```

...

You are running at the following UNICOS level
typhoon 8.0.3av roo.13 CRAY Y-MP

Here are the asynchronous product versions running on
this system at this time

```
as      version cc      version CRAY-T3D
4.0.3.4
CC      version 1.0.3.1
cdbx    version 8.1.0.6
cft77   version
debug   version 8.1.0.6
fmp      version 6.0.4.0
fpp      version 6.0
pascal  version 4.2.3.0
```


segldr version 8.0h

Cray Standard C Version 4.0.3.1 (097512) May 4 1995 23:23:41

Cray F90 Version 1.0 (1.68) 05/04/95 23:14:49

cf90: Cray CF90 Version 0.1.1.0 (037612) Thu May 4, 1995 23:14:50

Cray F90 Version 1.0 (1.73) 05/04/95 23:15:10

cf90: Cray CF90 Version 0.1.2.0 (118966) Thu May 4, 1995 23:15:12

PPLDR version 10.w - 01/02/95

typhoon[jpb]9164: uname -a

sn1049 typhoon 8.0.3av roo.13 CRAY Y-MP

sn1049 - Thu May 4 18:49:07 CDT 1995

B.2 Installed CRAY T3E Configuration

The following is the output from the `sysconf` command on the CRAY T3E. Note that although the system has 256 nodes only 80 of those nodes are contractor provided and thus eligible for benchmarking.

Hardware:

```
System serial number ..... SN6323
Mainframe type ..... CRAY-T3E
Mainframe subtype ..... T3EXXX
Number of available CPUs ..... 1
Cycle time in nanosecs (LPE 0x0f5) ... 3.3330
Clock ticks per second (LPE 0x0f5) ... 300000000
```

Software:

```
Operating system ..... UNICOS/MK
Release level ..... 1100
Kernel generation date and time ..... 02/03/97 10:08:00
Max PEs avail to application ..... 245
Level of Posix conformance ..... 199009
SECURE_SYS option ..... NORMAL
Operating system ticks per second .... 100
Posix Job Control implemented ..... YES
Max number of open files (current) ... 64
Max number of open files (limit) .... 64
Max value for User-ID ..... 60000
Max value for Process-ID ..... 100000
Max length of args for exec() ..... 49999
Max number of processes per user ..... 95
Max number of multi-group groups ..... 64
Max number of ptys ..... 128
Number of mount points configured .... 150
Number of users configured ..... 200
Number of I/O cache blocks ..... 7000
SCTrace enabled ..... NO
exec() saves IDs ..... YES
Running under a simulator ..... NO
Max size of timezone name ..... 128
Root privilege policy ..... ON
Posix privilege policy ..... OFF
TFMgmt priv policy under MLS ..... OFF
Enforce system high/low MAC ..... OFF
Secure mkdir(2) option ..... OFF
Size of kernel & tables (LPE 0x0f5) .. 10485760
User memory available (LPE 0x0f5) .... 123731968
```

The output of the `uname` command follows, with information on C and Fortran90 compiler version information.

```
sn6323 jim 1.3.160 unicosmk CRAY T3E  
Cray Standard C Version 5.0.3.0 (d29p35m275a35) Feb 28 1997 10:27:23  
Cray CF90 Version 2.0.3.1 02/28/97 10:27:20
```

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE May 1997	3. REPORT TYPE AND DATES COVERED Preliminary report	
4. TITLE AND SUBTITLE U.S. Army Engineer Waterways Experiment Station Major Shared Resource Center Benchmark Results and Analysis; Performance Level 1			5. FUNDING NUMBERS	
6. AUTHOR(S) John E. West, Alex R. Carrillo				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Engineer Waterways Experiment Station 3909 Halls Ferry Road Vicksburg, MS 39180-6199			8. PERFORMING ORGANIZATION REPORT NUMBER Technical Report ITL-97-3	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Corps of Engineers Washington, DC 20314-1000			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Available from National Technical Information Service, 5285 Port Royal Road, Springfield, VA 22161.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Upon installation of the Performance Level 1 target High-Performance Computing configuration in the Department of Defense U.S. Army Engineer Waterways Experiment Station Major Shared Resource Center, the Government's integration contract requires that the contractor demonstrate that the performance levels achieved by the installed configuration match those cited in the offeror's final pre-award proposal. This report details the analysis of the results of this demonstration, and enumerates several areas of concern that must be addressed before the witnessed benchmark is recommended as successfully complete.				
14. SUBJECT TERMS High-Performance Computing Systems Analysis Operations research Computer Benchmarks			15. NUMBER OF PAGES 41	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	